

Remarks

Entrance of this amendment and reconsideration of the pending claims are respectfully requested. Upon entrance of this amendment, claims 12-19 & 31-38 will be pending.

By this paper, new system claims 31-38 are added for the Examiner's consideration. These claims principally correspond to amended method claims 12-19 also under consideration. In addition, by this paper, applicants amend independent claim 12 to more particularly characterize the collective processing of data by the dedicated collective offload engine in applicants' invention as being without use of any software tree. Additionally, applicants amend claim 12 to specify that the result produced by the dedicated collective offload engine is produced in a deterministic time based on the collective processing. Support for these amendments can be found throughout the application as filed. For example, reference specification paragraphs [0015], [0016] & [0028]. Thus, no new matter is believed added to the application by any amendment presented.

In the Office Action, prior claims 12-19 were rejected under 35 U.S.C. §103(a) as being unpatentable over Burianek et al. (U.S. Patent No. 7,082,457; hereinafter Burianek) in view of Bernardo (U.S. Patent No. 6,766,517; hereinafter Bernardo). This rejection is respectfully traversed, and reconsideration thereof is requested.

Applicants' independent claims recite providing, by a *dedicated collective offload engine* coupled to a switch fabric in a distributed parallel computing system, collective processing of data. There is no dedicated collective offload engine in Burianek as the term is employed in Applicants' specification and claims. Further, there is no collective processing of data in Burianek. In the Office Action, Applicants' recited dedicated collective offload engine is analogized to server 215 of Burianek. This analogy is respectfully traversed.

Server 215 in Burianek is described as a project management central server that directs signals sent to and from the components of the distributed computing environment. This server includes a delegation component which sends and receives information about project tasks stored in the database 210. Thus, server 215 in Burianek is a conventional server system. This is distinguished from Applicants' *dedicated collective offload engine*, which provides collective processing of data. In Applicants' invention, the *dedicated collective offload engine is a*

hardware device coupled to the switch fabric and further, applicants recite that the collective processing of data by the dedicated collective offload engine *is without use of any software tree*. In Burianek, the processing described is implemented in software, while in Applicants' recited invention, collective processing is implemented in a hardware device, that is, the dedicated collective offload engine without use of any software tree. The characterization *without use of any software tree* is clearly not taught or suggested by Burianek, which describes typical software program implemented processing. *A software program by its very nature necessarily includes or is a software tree*. In contrast, applicants' invention is implemented without use of any software tree, that is without use of a conventional software program. This is because applicants do not rely on a processor-based implementation for their dedicated collective offload engine. As noted above, in applicants' invention the dedicated collective offload engine is a hardware device and the collective processing is achieved without any software tree.

Applicants recite a data processing method which includes collective processing of data from the at least some processing nodes of the multiple processing nodes of a distributed, parallel computing system. The collective processing implements a collective operation on the data from the at least some processing nodes. The phrases "*collective processing*" and "*collective operation*" are terms of art which *refer to a particular type of data processing*. A collective operation is conventionally an arithmetic operation executed across data from multiple nodes of a distributed, parallel computing system, with results being provided to multiple nodes. Thus, a collective operation is an n:n operation.

As explained in Applicants' "Background of the Invention", implementation of collective processing typically includes using a software tree approach, wherein message passing facilities are used to form a virtual tree of processes. A drawback to this approach is the serialization of delays at each stage of the tree. These delays are additive in the overall overhead associated with the collective processing. Furthermore, the software tree approach results in a theoretical logarithmic scaling latency of the overall collective processing versus system size. Due to interference from daemons, interrupts and other background activity, cross traffic, and the unsynchronized nature of independent operating system images and their dispatch cycles, measured values of scaling latency are usually significantly worse than theoretical values. Responsive to this issue, Applicants describe a novel *collective processing approach* with mitigates the large latency associated with the software tree implementation. In Applicants'

approach, a dedicated collective offload engine, which is a hardware device coupled to the switch fabric, is employed to provide the collective processing of data from the multiple processing nodes. Applicants' hardware device is a specialized device dedicated to providing the collective processing in hardware of the data, and the collective processing implements a collective operation on the data. The collective processing occurs in hardware since the device itself is a hardware device, and no software tree (or software program) is employed to perform the collective operation.

The above-noted novel aspects of applicants' invention are further reinforced in amended claim 12 and new claim 31 wherein applicants recite *producing by the dedicated collective offload engine a result in deterministic time based on the collective processing*. Support for this amendment can be found throughout the application as filed. For example reference specification paragraph [0028], where applicants describe that the collective processing presented promotes deterministic performance. Since performance is time based, applicants' processing produces a result in deterministic time. This is distinct from the software implemented approaches described in Burianek and Bernardo. In both patents, the computer or server comprises a processor which executes a software program that is susceptible to delays, such as interrupts in the processing requested. Applicants' dedicated collective offload engine advantageously eliminates the prior art's processor based collective processing wherein the result is indeterministic in time.

In applicants' invention, the dedicated collective offload engine is a hardware device that is coupled to the switch fabric and which communicates with the at least some processing nodes of the distributed parallel computing system across the switch fabric. Applicants' dedicated collective offload engine advantageously collectively processes in hardware, without a software tree, received data from the at least some processing nodes, and provides a result in a deterministic time based on the collective processing.

Still further, applicants' independent claims recite that the dedicated collective offload engine is a hardware device that is a specialized device dedicated to provide the collective processing of data received from at least some processing nodes. The servers and computer systems described in the applied art are not dedicated devices *per se*. In both cases, general purpose processors are employed to provide a variety of processing options and functions.

Applicants further specify that this specialized hardware device includes *a dispatcher built from field programmable gate arrays*, and a pipelined arithmetic logic unit, wherein the dispatcher controls collective processing of the received data by the arithmetic logic unit. Applicants respectfully submit that the Office Action does not address applicants' above-noted characterizations in the independent claims presented. *Specifically, the applied art, and the Office Action do not describe a dispatcher built from field programmable gate arrays per se, let alone a dispatcher built from field programmable gate arrays that is part of a specialized hardware device as applicants recite in the independent claims presented.*

For at least the above-noted reasons, applicants respectfully request allowance of the independent claims presented herewith. The dependent claims are believed allowable for the same reasons as the independent claims, as well as for their own additional characterizations. For example, dependent claims 18 & 36 recite a plurality of cascaded, dedicated collective offload engines, which are connected to communicate with the at least some processing nodes across a switch fabric and which together provide the collective processing data from the at least some processing nodes. No similar teaching is believed provided in the applied and known art.

All claims are believed to be in condition for allowance, and such action is respectfully requested.

Should the Examiner have reservations regarding the patentability of any claim(s) presented, Applicants' undersigned representative respectfully requests the opportunity for an Examiner Interview to discuss the claim(s) in the hope of advancing prosecution of the subject application.

Respectfully submitted,



Kevin P. Radigan, Esq.
Attorney for Applicants
Registration No.: 31,789

Dated: February 16, 2009.

HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 Columbia Circle
Albany, New York 12203-5160
Telephone: (518) 452-5600
Facsimile: (518) 452-5579